

LISTING OF CLAIMS

1. (Previously Presented) A method comprising:
maintaining a synchronous copy of a data change log at a primary node, wherein
said data change log at said primary node is associated with a primary data
volume of said primary node; and
said synchronous copy of said data change log is maintained at a data recovery
node;
asynchronously updating a secondary data volume of a secondary node using said
synchronous copy of said data change log.
2. (Previously Presented) The method of claim 1, wherein said secondary data volume is
updated in response to detecting a failure of the primary data volume.
3. (Cancelled)
4. (Previously Presented) The method of claim 1, wherein said maintaining said
synchronous copy comprises:
receiving a request to perform a write operation on said primary data volume;
storing data associated with said write operation substantially simultaneously on said data
change log and said synchronous copy of said data change log in response to said
receiving.
5. (Cancelled)
6. (Cancelled)
7. (Previously Presented) The method of claim 2, wherein said updating comprises:
receiving a manual update initiation indication; and
updating said secondary data volume using said real-time copy of said data change log in
response to said receiving.

8. (Previously Presented) The method of claim 1, wherein said synchronous copy of said data change log comprises a plurality of entries; and said updating comprises:
 - identifying an entry of said plurality of entries as corresponding to an incomplete write operation on said primary data volume; and
 - updating said secondary data volume using said entry.
9. (Previously Presented) The method of claim 1, wherein said updating comprises:
 - copying a block of data from said synchronous copy of said data change log to a staging log at said secondary node, said block of data comprising a plurality of entries;
 - applying each of said plurality of entries to a data change log at said secondary node in response to said copying; and
 - updating said secondary data volume using said data change log at said secondary node.
10. (Previously Presented) The method of claim 2, further comprising:
 - detecting a recovery of said primary data volume; and
 - resynchronizing said primary data volume and said secondary data volume in response to said detecting.
11. (Previously Presented) A machine-readable medium comprising a plurality of instructions, wherein said plurality of instructions when executed implement a method comprising:
 - maintaining a synchronous copy of a data change log at a primary node, wherein said data change log at said primary node is associated with a primary data volume of said primary node; and
 - said synchronous copy of said data change log is maintained at a data recovery node;
 - asynchronously updating a secondary data volume of a secondary node using said synchronous copy.
12. (Previously Presented) The machine-readable medium of claim 11, wherein said secondary data volume is updated in response to detecting a failure of the primary data volume.
13. (Cancelled)

14. (Previously Presented) The machine-readable medium of claim 11, wherein said maintaining said synchronous copy comprises:
- receiving a request to perform a write operation on said primary data volume;
 - storing data associated with said write operation substantially simultaneously on said data change log and said synchronous copy of said data change log in response to said receiving.
15. (Cancelled)
16. (Cancelled)
17. (Previously Presented) The machine-readable medium of claim 11, wherein said synchronous copy of said data change log comprises a plurality of entries; and said updating comprises:
- identifying an entry of said plurality of entries as corresponding to an incomplete write operation on said primary data volume; and
 - updating said secondary data volume using said entry.
18. (Previously Presented) The machine-readable medium of claim 11, wherein said updating comprises:
- copying a block of data from said synchronous copy of said data change log to a staging log at said secondary node, wherein said block of data comprises a plurality of entries;
 - applying each of said plurality of entries to a data change log at said secondary node in response to said copying; and
 - updating said secondary data volume using said data change log at said secondary node.
19. (Previously Presented) A data processing system comprising:
- means for maintaining a synchronous copy of a data change log at a primary node, wherein
 - said data change log at said primary node is associated with a primary data volume of said primary node; and
 - said synchronous copy of said data change log is maintained at a data recovery node;

means for asynchronously updating a secondary data volume of a secondary node using said synchronous copy of said data change log.

20. (Previously Presented) The data processing system of claim 19, wherein secondary data volume is updated in response to detecting a failure of the primary data volume.

21. (Cancelled)

22. (Previously Presented) The data processing system of claim 19, wherein said means for maintaining said synchronous copy comprises:

means for storing data associated with a requested write operation on said primary data volume substantially simultaneously on said data change log and said real-time copy of said data change log.

23. (Cancelled)

24. (Cancelled)

25. (Previously Presented) The data processing system of claim 19, wherein said synchronous copy of said data change log comprises a plurality of entries; and said means for updating comprises:

means for identifying an entry of said plurality of entries as corresponding to an incomplete write operation on said primary data volume; and
means for updating said secondary data volume using said entry.

26. (Previously Presented) The data processing system of claim 19, wherein said means for updating comprises:

means for copying a block of data from said synchronous copy of said data change log to a staging log at said secondary node, said block of data comprising a plurality of entries;

means for applying each of said plurality of entries from said staging log to a data change log at said secondary node; and

means for updating said secondary data volume using said data change log at said secondary node.

27. (Previously Presented) A data processing system comprising:
a storage element to store a synchronous copy of a data change log at a primary node,
wherein said data change log at said primary node is associated with a primary
data volume of said primary node; and
a recovery module configured to asynchronously update a secondary data volume of a
secondary node using said synchronous copy of said data change log.
28. (Previously Presented) The data processing system of claim 27, wherein said recovery
module updates the secondary data volume in response to a failure of said primary data volume.
29. (Previously Presented) The data processing system of claim 27, further comprising:
a volume management module configured to mirror data to be written to said data change
log to said synchronous copy of said data change log.
30. (Cancelled)
31. (Previously Presented) The data processing system of claim 27, wherein
said synchronous copy of said data change log comprises a plurality of entries; and
said recovery module comprises a failover management module configured to identify an
entry of said plurality of entries as corresponding to an incomplete write operation
on said primary data volume and update said secondary data volume using said
entry.
32. (Previously Presented) The data processing system of claim 27, wherein said recovery
module comprises:
a failover management module configured to copy a block of data comprising a plurality
of entries from said real-time copy of said data change log to a staging log at said
secondary node.
33. (Cancelled)
34. (Cancelled)

35. (Previously Presented) A method comprising:
maintaining a synchronous copy of a data change log at a primary node, wherein
said data change log at said primary node is associated with a primary data
volume of said primary node; and
said synchronous copy of said data change log is maintained at a data recovery
node; and
asynchronously replicating data to be written to said primary data volume from said
primary node to said secondary node.
36. (Cancelled)
37. (Previously Presented) The method of claim 35, further comprising:
detecting a failure of said primary data volume; and
updating a secondary data volume of said secondary node using said synchronous copy of
said data change log in response to said detecting.